

**A more recent version of this article has been accepted for publication in:
Energy Economics, 2015**

Forecasting Gasoline Prices in the Presence of Edgeworth Price Cycles

Michael D. Noel* and Lanlan Chu†

May 27, 2015

Abstract

Forecasting is a central theme in economics. The ability to forecast prices enables economic agents to make optimal decisions for the present and future. In this article, we investigate if and how gasoline prices can be forecast in retail gasoline markets that are subject to high-frequency, asymmetric price cycles known as Edgeworth price cycles. We examine a series of purchase timing decision rules and a series of feasible forecasting algorithms for updating those rules over time. We find that, in the presence of cycles, agents in our five Australian markets can systematically reduce purchase prices below market average the equivalent of 11 to 15 U.S. cents per gallon, using simple decision rules and feasible forecasting algorithms.

1 Introduction

Forecasting prices is one of the central themes in economics. The ability to predict future prices and other outcomes based on past information is important in helping economic agents make optimal intertemporal decisions. One industry in which the ability to predict prices day-to-day can potentially be important to consumers is retail gasoline. Gasoline is purchased frequently and there is ample scope for adjusting purchases by a few days if expected prices are different than current prices. Since a significant portion of many consumers' disposable incomes is spent on gasoline, and aggregate short run elasticities are close to zero (Hughes et al. (2008)), changes in gasoline prices can produce large income effects (Gicheva et al. (2010)) and there is potential gain in trying to predict them.

Obviously, if gasoline prices change little day to day, the gain from forecasting prices is limited. Agents typically purchase every week or two and it is rarely worthwhile for agents to stockpile more

*Corresponding author. Michael D. Noel, Associate Professor, Department of Economics, Texas Tech University, Lubbock, Texas, michael.noel@ttu.edu.

†Lanlan Chu, Department of Economics, Texas Tech University, Lubbock, Texas, lanlan.chu@ttu.edu.

than a tankful of gasoline. However, when day to day volatility is very high, the potential for gains from forecasting is also potentially high.

In this article, we examine the day-to-day forecastability of gasoline prices in markets that experience significant price volatility in the form of high-frequency, asymmetric retail price cycles known as Edgeworth price cycles. Figure 1 shows an example of Edgeworth price cycles in Melbourne, Australia, from July to December, 2010. Prices exhibit a high-frequency asymmetric retail price cycle in which prices rise over the course of a few days and then fall over the course of the next week or two. Similar cyclical patterns are found in gasoline prices in all five major Australian cities, as well as in parts of Canada, the United States, and various European countries. In the Melbourne figure, the change in prices from the bottom of the cycle to the top ranges from 8.7 to 15.3 Australian cents per liter (31.2 to 55.0 U.S. cents per gallon) across the figure.¹ In other words, purchasing at the peak would cost up to 55 U.S. cents per gallon more (or about \$9 more on a fillup) than purchasing at the trough just one or two days earlier. The ability to forecast the troughs with some sort of timing strategy can potentially avoid this.²

< Figure 1 placed about here >

To our knowledge, the only study that has looked at the predictability of troughs along a cycle is Noel (2012), who examined prices in Toronto, Canada, in 2001. Noel (2012) found that optimized timing strategies, if used, would have reduced prices paid 2 – 4% over that sample, depending on the complexity of the strategy.

There is a key limitation of that study, however. The sample period is short, just over four months for a single city, and includes only 18 complete cycles. For power, the author must rely on the full sample to estimate optimal purchase timing decision rules and the price reductions from using those rules. But it means the reported decision rules are quite special. They are those that would have worked best if agents had *perfect foresight* and were using them from the start

¹One cent per liter in Australian dollars is equivalent to 3.6 US cents per gallon, at the average exchange rate over the sample.

²Rising or falling crude oil prices can help forecast retail prices due to the transmission lag (see Noel and Lewis (2011)), but the gains are limited by the relatively slow pace of change compared with the typical frequency of purchases. Also, since the short run volatility in retail prices in the presence of Edgeworth price cycles are not reflected in futures markets for crude or wholesale gasoline, agents cannot rely on those futures markets to predict the timing of the cycle troughs.

of the sample, even though they are not known until the sample is analyzed ex post. Similarly, the reported price reductions are special in that they depend on perfect foresight in-sample and are only meaningful out-of-sample if cycle characteristics are perfectly stable over time.³ This is generally not the case. Noel (2012) notes this and warns that changes in cycle conditions change optimal choices and thus decay the estimated price reductions over time. Updating decision rules is necessary.

But how to update or, more generally, how to forecast prices and decision rules forward is a different question and not one that can be addressed in that study. Nor can it be determined if price reductions are possible at all with only ex-ante forecasts and without perfect foresight. In fact, many important questions remain.

Can purchase prices be systematically lowered using feasible forecasting algorithms in the absence of perfect foresight? If so, what purchase timing decision rules and classes of decision rules reduce prices the most? (A class of purchase timing decision rules is a collection of rules that all share a common trigger mechanism and differ only in the specific threshold values that trigger purchasing.) How does the effectiveness of decision rules within a class change over time with changes in the shapes of the cycles? Are some classes of rules more stable than others? In other words, which classes of rules tend to "hold their value" better and which tend to have optimal triggers that "bounce around" a lot making them difficult to predict one period to the next. (Since cycle characteristics change in different ways, and different classes are based on different characteristics, classes can vary in stability and long term effectiveness.)

In terms of how to update forecasts over time, ruling out perfect foresight, what is the most effective algorithm for forecasting the next set of optimal decision rules on the basis of prior and known data? How do the price reductions under these forecasts compare to those under perfect foresight? What is the best combination of decision rule classes and forecasting algorithms to use?

Our data is sufficiently rich that we are able to address all these questions. In contrast to 18 complete cycles, our data includes retail and wholesale prices for five cities in Australia over six and a half years for a total of 1,340 complete cycles. There is also significant variation in the cycles

³Absent stable cycles, the "average" optimal rule calculated over the T period sample no longer reflects the optimal rule for period $T + s$, $s = 1, 2, 3...$ (or even the optimal rule in some periods $t < T$).

across cities and over time.

Noel (2012) credibly made the point that cycle troughs are in principle predictable. The contribution of this paper is to operationalize that idea and determine if, and if so how, cycle troughs can be forecast feasibly, with only prior and known information, in the absence of perfect foresight, and in the face of constantly evolving cycles.

To preview results, we indeed find that perfect foresight, the best case scenario used by Noel (2012), overstates the price reductions obtainable by agents. The amount varies by class. However, we find that *feasible* forecasting algorithms, relying on only prior and known data, still produce significant and large reductions in price for certain classes. We conclude that cycles are meaningfully predictable and price reductions are obtainable, not only in principle, but also in practice.

We finish with some policy implications of our work. The concern surrounding short run volatility in gasoline prices has led to numerous investigations by competition authorities around the world and, in some cases, direct policy interventions. Legislative responses range from mandating an online price census that consumers can access (e.g. Greece, Spain, France, Cyprus, Portugal) to mandatory price increase pre-notification laws (e.g. Australia, Austria) to outright price regulation (e.g. eastern Canada). While the ability to forecast periods of high and low gasoline prices has important welfare implications, relatively little attention has been given to understanding the extent to which prices can be forecast. In this article, we examine the day-to-day forecastability of gasoline prices in markets that exhibit Edgeworth price cycles. We discuss how the forecasting function could be performed by consumers, but more feasibly, by a federal agency or private firm, with simple purchase timing rules based on those forecasts passed down to consumers for practical use.

2 Literature

Edgeworth Price Cycles are a competitive pricing phenomenon in which prices cycle asymmetrically at a high frequency. They were first conceptualized by Edgeworth (1925) and theoretically formalized by Maskin & Tirole (1988). The model has been extended to asymmetric firms by Eckert (2003) and to a wider range of different competitive assumptions, such as the inclusion of stochastic marginal costs and varying market structures, by Noel (2008).

The model is as follows. Two infinitely-lived profit-maximizing firms compete in a homogeneous Bertrand pricing game by setting prices in an alternating fashion – one firm sets its price in even periods and the other in odd periods – and once set, the price for that firm is fixed for two periods. If Firm 1 adjusts its price in period t , then $p_t^1 = p_{t+1}^1$ and $p_t^2 = p_{t-1}^2$. Prices are chosen from a discrete price grid of L possible prices. Marginal cost c is constant and fixed costs are zero. Each firm earns current period profits of:

$$\pi_t^i(p_t^1, p_t^2, c_t) = D^i(p_t^1, p_t^2) * (p_t^i - c) \quad (1)$$

where D^i is the standard Bertrand formulation.

The strategies of each firm are allowed to depend only on the payoff-relevant state at t . That is, they depend only on the price set by the other firm in the previous period, current demand, and current marginal cost. The Markov Perfect equilibrium strategies are given by R^1, R^2 , where $(p_t^1)^* = R^1(p_{t-1}^2)$, $(p_t^2)^* = R^2(p_{t-1}^1)$ and p_{t-1}^j is the price chosen by firm j in period $t - 1$ which remains in effect in period t .

Let $V^1(p_{t-1}^2)$ be the Firm 1's value function when Firm 2 adjusted its price to p_{t-1}^2 in the previous period, and Firm 1 adjusts its price in the current period. Let $W^1(p_{s-1}^1)$ be Firm 1's value function when it has set price p_{s-1}^1 in the previous period and Firm 2 is about to adjust its price. V^1 and W^1 are written:

$$V^1(p_{t-1}^2) = \max_{p_t} [\pi_t^1(p_t, p_{t-1}^2, c) + \delta W^1(p_t)] \quad (2)$$

$$W^1(p_{s-1}^1) = \mathbb{E}_{p_s} [\pi_s^1(p_{s-1}^1, p_s, c) + \delta V^1(p_s)] \quad (3)$$

and similar equations are found for V^2 and W^2 . The discount factor is δ . The expectation shown in W^1 is taken with respect to the distribution of R^2 .

The system is in equilibrium when V^i and W^i are independent of t , i.e. they are fixed point L -vectors. Maskin and Tirole (1988) show that both focal price equilibria and Edgeworth Price Cycle equilibria are possible in this setting. The Edgeworth Price Cycle equilibrium takes the form:

$$R^i(p^j) = \begin{cases} \bar{p} & \text{for } p^j > \bar{p} \\ p^j - k & \text{for } \bar{p} \geq p^j > \underline{p} \\ c & \text{for } \underline{p} \geq p^j > c \\ c & \text{with probability. } \mu(\delta) \text{ for } p^j = c \\ \bar{p} + k & \text{with probability. } 1 - \mu(\delta) \text{ for } p^j = c \\ c & p^j < c \end{cases} \quad (4)$$

< Figure 2 placed about here >

Figure 2 shows an example of equilibrium price paths with $c = 0$. The practical mechanics are straightforward. Firms undercut one another until the price gets relatively close to marginal cost, at which point one firm drops its price immediately to marginal cost. Firms then play a war of attrition, each playing a mixing strategy, mixing between raising the price back to $\bar{p} + k$ and maintaining price equal to marginal cost. By holding at marginal cost, the firm creates the possibility that the other firm will raise price first, enabling the first firm to capture sales at the highest possible price, \bar{p} , on its next turn. (Only the lowest price firm in a period makes sales.) The risk is that the other firm will hold fast at cost too, meaning the first firm will have foregone a period of profit-taking only to find itself in the same situation again. Eventually, one firm does raise price to $\bar{p} + k$, the other undercuts, and tit-for-tat undercutting begins again.

An interesting feature of the cycles, and one that is important for us, is that they should have predictable qualities. Of the six possible price responses listed above, all but the last are observed on the equilibrium path and in a systematic way. Only the mixing adds some variance to the periodicity.

The situation is more complex in the Noel (2008) extension of the model. There is variation due to stochastic marginal costs and changing demands, market structures and other considerations. The best response functions depend on the realization of c_t , so $(p_t^1)^* = R^1(p_{t-1}^2, c_t)$, $(p_t^2)^* = R^2(p_{t-1}^1, c_t)$ and functions V and W are similar to the above, but with a different D^i , and written as expectations over the distribution of c . The system is solved numerically and the best response functions are similar to the above but with all six possible responses observed on the equilibrium

path.⁴ Variation in costs and demands, and difficulty in completing market-wide price increases imply less stable cycle characteristics over time and more complexity in forecasting troughs.

A large empirical literature has developed over the past fifteen years that examines high frequency asymmetric cycles in retail gasoline markets. Cycles have been observed in the United States, Canada, Australia, and various European countries. Castanias and Johnson (1993) first noted the similarity of the asymmetric retail gasoline price cycles in certain U.S. cities in the 1960s and 1970s with the then new model of Maskin & Tirole (1988). Then beginning in the 2000s, a series of papers tested the empirical cycles against predictions of the Eckert (2003) and Noel (2008) extensions of the Edgeworth Price Cycle model. These include Eckert (2003), Noel (2007a), Noel (2007b), and Atkinson (2009) for Canada, Lewis (2009), Doyle et al. (2010), and Lewis (2012) for the U.S., and Wang (2009a) and Wang (2009b) for Australia. They generally find support for the Edgeworth Price Cycles model. Other papers including De Roos and Katayama (2013) for Australia and Foros and Steen (2013) for Norway find results consistent with Edgeworth Price Cycles but remain agnostic as to the underlying process.⁵

Several papers look at the effect of Edgeworth Price Cycles on price levels and found that cycles do not increase prices, contrary to popular perception. Papers include Noel (2002) and Noel (2015) for Canada, and Doyle et al. (2010) and Zimmerman et al. (2013) for the U.S. Other papers have looked at the more subtle interaction between Edgeworth Cycles and the so-called Rockets and Feathers phenomenon (Eckert (2002), Noel (2009)), or how the cycles can improve market recovery following cost shocks (Lewis (2009), Lewis and Noel (2011)). Noel (2012) finds that the ability to time cycle troughs can result in a further decrease in price levels with cycles.

3 Data and Methodology

In this paper, we exploit a more comprehensive dataset to examine the forecastability of cycle troughs in the five major Australian capital cities – Sydney, Melbourne, Brisbane, Adelaide, and Perth. We use daily average retail prices for each city for regular grade gasoline, obtained from Informed Sources, and daily average terminal gate prices (the wholesale price) for regular grade

⁴Cost changes can cause negative margins in a given period.

⁵The results of this article do not depend on the underlying process.

gasoline, obtained from Orima Research. Prices are expressed in Australian cents per liter. The sample period runs from January 1, 2007 to June 30, 2013.

There are five stages to the analysis. The first empirically describes the anatomy of the cycle. We estimate the magnitude of price changes during relenting phases and undercutting phases and the rate of switching between them. We then use these estimates to construct measures of the structural characteristics of the cycle, i.e. its period, amplitude and asymmetry, following Noel (2007a). We also discuss the variance in cycle characteristics, which has implications for the stability of purchase timing decision rules over time.

In the second stage, we identify five classes of purchase decision rules that potentially hold predictive power. One is new to the literature. A class of decision rules is defined as a collection of rules that use a similar concept and logic, where each rule within a class is differentiated only by having a different trigger point. For example, the class of "Spike and Wait" decision rules call for purchasing X days after the last peak. Each decision rule within the class is differentiated only by a different X . We perform a series of preliminary regressions to test whether each class has potential value in predicting cycle troughs.

The third stage begins the main analysis. First, we simulate intertemporal purchase profiles for agents under each decision rule within each class and over a given interval of time (generally a half-year). We calculate the price reduction from using each specific decision rule within a class, relative to myopic, or random, purchasing. This yields curves relating price reductions to an ordered set of decision rules in each class, i.e. reduction profiles. We then find the optimal decision rule which maximizes the price reduction (or equivalently, maximizes the reduction profile) in each class over each interval. While this optimal rule is not known by the agent until after the interval is over, it can be used to help forecast future maxima.

In the fourth stage, we take a first look at decision rule stability. The less variation in the reduction profile maxima from one interval to the next in a given class, the more stable will be the class. Also, the weaker is the concavity in the reduction profile for a given class and interval, the more stable the class. The idea is that more stable maxima over time are easier to forecast and less concave reduction profiles do not punish as severely for missed forecasts.

In the fifth and final stage, we forecast future prices based on only prior and known data and

test whether price reductions are feasible in the absence of perfect foresight. We consider a variety of forecasting algorithms based on first-order and higher-order Markov processes. We compare the prices actually obtained under each forecasting algorithm with those that would be obtained under the case of no forecasting at all (myopic purchasing) and the case of perfect foresight.

4 Results

4.1 Anatomy of a Price Cycle

We estimate the magnitude of price changes and the frequency of phase switching and, from these estimates, derive measures of the structural characteristics of the cycle. The effectiveness of decision rule classes lies in their ability to predict the cycle troughs by predicting the magnitude of various cycle characteristics.

We define a relenting phase as a sequence of consecutive price increases and the undercutting phase as a sequence of consecutive price decreases in between relenting phases.⁶

Define $\Delta RETAIL_{mt}^i = RETAIL_{mt} - RETAIL_{m,t-1}$ as the change in the retail price in market m at time t conditional on phase i . We assume retail prices change according to the function:

$$\Delta RETAIL_{mt}^i = X_{mt}^i \beta^i + \varepsilon_{mt}^i \quad i = \{R, U\} \quad (5)$$

where $\varepsilon_{mt}^i \sim N(0, \sigma_i^2)$ and X_{mt} is a set of explanatory variables. Let the predicted values of $\Delta RETAIL_{mt}^i$ from this regression be given by $\alpha_{mt}^i(X_{mt}^i)$.⁷

Let I_{st} equal R or U when the market is in a relenting or undercutting phase respectively. Then the probability that a market m at time t transitions from phase i to phase j is given by the probit

⁶For relenting phases, we require the cumulative price increase (amplitude) to be at least $h = 2$ cents per liter. Other reasonable values for h , including zero, yields similar results. The purpose of a non-zero h is that, rarely, there is a small isolated average price increase (of usually less than a cent, instead of the usual cumulative increase of 8–11 cpl), generally near a trough, in what is otherwise a clear undercutting phase. These are due to "false starts", i.e. failed attempts to start a new relenting phase (see Noel (2008)). A small non-zero h avoids miscategorizing these as full relenting phases in the rare instances it does occur.

⁷We do not reject the null hypothesis of stationarity of the $\Delta RETAIL$ series. We do reject the null hypothesis of stationarity in the $RETAIL$ and TGP series in levels individually, but not surprisingly find they are cointegrated, consistent with previous work.

specification:

$$E(I_{mt} = j \mid I_{m,t-1} = i) = \Phi[W_{mt}\theta], \quad i, j = \{R, U\} \quad (6)$$

where W is a matrix of explanatory variables. Let the predicted values of this equation be given by $\lambda_{mt}^{ij}(W_{mt}^i)$. By the adding up constraint, $\lambda_{mt}^{iU} = 1 - \lambda_{mt}^{iR}$.

In the preliminary specification, the X_{mt}^U , X_{mt}^R , and W_{mt}^U are all vectors of ones, yielding constant values for λ^{ij} and α^i and average measures of cycle characteristics for each market. A separate regression for each market is freely estimated.

The expected duration of a particular phase i is then calculated by:

$$E(DURATION_m^i) = (1 - \lambda^{ii})^{-1} \quad (7)$$

The expected period of the cycle is the sum of the expected length of a relenting phase and the expected length of an undercutting phase, calculated as:

$$E(PERIOD_m) = (1 - \lambda^{RR})^{-1} + (1 - \lambda^{UU})^{-1} \quad (8)$$

Cycle amplitude can be calculated in one of two ways:

$$\begin{aligned} E(AMPLITUDE_m) &= \alpha^R(1 - \lambda^{RR})^{-1} \\ E(AMPLITUDE_m) &= -\alpha^U(1 - \lambda^{UU})^{-1} \end{aligned} \quad (9)$$

and cycle asymmetry can also be calculated in one of two ways:

$$\begin{aligned} E(ASYMMETRY_m) &= -\alpha^R/\alpha^U \\ E(ASYMMETRY_m) &= (1 - \lambda^{RR})^{-1}(1 - \lambda^{UU}) \end{aligned} \quad (10)$$

Given that prices are similar at the start and end points of the sample, each pair of calculations roughly agree.⁸

⁸We report only the former of each set in our results.

< Table 1 placed about here >

< Table 2 placed about here >

Table 1 shows the mean within-phase price changes and the phase transition probabilities from R to U for the five markets. Table 2 shows the derived cycle characteristics. As an example, the mean average price increase in a relenting phase in Melbourne is 3.82 cpl (α^R) and the mean duration of the relenting phase is $(1 - 0.606)^{-1} = 2.54$ days (from the time the first stations increase price to when virtually all remaining stations have). This yields an average amplitude of 9.70 cpl. The mean period of the cycle is 8.68 days, including 6.14 days of undercutting on average, and asymmetry is measured at 2.42 (unit-free).

Comparing across cities, mean cycle periods are similar, ranging from 8.1 days in Perth to 8.7 days in Adelaide, averaged across the whole sample. Mean relenting phase durations are around 2.5 days, typically the result of a small handful of stations raising price on the first day and the wide majority raising price on the second day, with sometimes a few holding out to the third day. Mean cycle asymmetry is about 2.5. Amplitudes are more varied across cities, ranging from 7.02 cents per liter (cpl) in Perth to 11.65 cpl in Adelaide, on average.

The standard errors (of the means) show some variance in our mean estimates, but more relevant for forecasting purposes are the standard deviations around the mean value for individual cycles. Given $N = 272$ full cycles in a city, standard deviations are about sixteen times the standard error of the means. That means there is a large amount of variance in the individual cycle characteristics over time. Since decision rules are based on cycle characteristics, instability in cycle characteristics can lead to instability of decision rule classes over time.

4.2 Classes of Decision Rules

We identify five classes of purchase timing decision rules. All classes work by attempting to predict when prices will be at the trough and call for purchasing just in advance of it. If there is regularity in the cycle period, decision rules can be based on predicting which day to buy. We refer to these as horizontal rules. If there is regularity in the cycle amplitude, decision rules can be based on predicting at which price (or firm margin) to buy. We refer to these as vertical rules. A third type

of rule, neither horizontal nor vertical, focuses on regular anomalies in the distribution of prices during relenting phases, and we call this a distributional rule.

Position Based Decision Rules

We first address two classes of vertical rules. The first is the class of Position Based decision rules. Under these rules, at every time t the agent compares the wholesale price (terminal gate price) TGP_{mt} and the retail price $RETAIL_{mt}$ and calculates the current "position" of the cycle, $POSITION_{mt} = RETAIL_{mt} - TGP_{mt}$.⁹ Terminal gate prices are posted publicly each day. The class calls for purchase when the value of $POSITION$ drops below a given critical level μ_{PB} . Rules within the class are differentiated only in that they trigger at different values of μ_{PB} .

We need a contingency rule for what occurs when the value for μ_{PB} is set so low that it is not reached at the troughs, reached only occasionally, or is reached regularly but the troughs are still too far apart for a reasonable purchase frequency. We handle this by assuming if a μ_{PB} is not reached within $w = 12$ days since the previous purchase, the agent must purchase regardless of its current value or the current price. We make equivalent assumptions for all other classes as well.¹⁰

Spike and Drop Decision Rules

The second class of vertical rules is the Spike and Drop Class. This is new to the literature. Under these rules, at every time t , the agent calculates the cumulative decrease in price since the previous peak. The class calls for purchasing when the cumulative price decrease, $DECREASE_{mt} = \sum_{s=0}^{DAY_{S_{mt}}} \Delta RETAIL_{mt-s}$ exceeds a given critical level μ_{SD} . The previous peak price is defined as $p_{m,PEAK} = \{p_{mt} \mid I_{mt} = R, I_{m,t+1} = U\}$ and $DAY_{S_{mt}}$ is the number of days since the peak. Rules within the class are differentiated only in that they trigger at different values of μ_{SD} .

Spike and Wait Decision Rules

⁹We refer to this as cycle $POSITION$ instead of cycle margin, consistent with the literature, since consumers do not care about margins directly but rather about the position of prices relative to the estimated trough. Also, since retail minus rack is a proxy for margin and our results do not depend on the quality of this proxy, $POSITION$ is a more appropriate term in this setting.

¹⁰The value for w is based on the average time between fill-ups for Australian drivers. The average driver drives 15,500 kilometers per year, at a fuel consumption rate of 11.3 liters per 100 km, and a tank capacity of 60 liters. Our results are robust to different assumptions about w , with price reductions a little higher for higher w drivers (more flexible) and a little lower for lower w drivers (less flexible).

The next two classes consist of horizontal rules. The class of Spike and Wait decision rules is based on serial correlation in the periodicity of cycles. The class calls for purchasing when the number of days since the last peak, $DAY S_{mt}$, reaches a critical level, μ_{SW} .

Calendar Based Decision Rules

The class of Calendar Based decision rules is based on the idea that relenting phases can be synchronized with specific events on the calendar. There is evidence that relenting phases in Australian markets were concentrated in some years on certain days of the week. Let $DOW_t^1 \dots DOW_t^7$ be indicator variables for each day of the week. Rules that call for purchasing on a specific day of the week, i.e. when DOW_t^i reaches the chosen value μ_{CB} .

Spike and Buy Decision Rule

The final class of decision rules is the Spike and Buy Class. It is neither horizontal nor vertical, but distributional. It is unique in that it is a singleton class - there is only one decision rule in it. It is based on the observation that marketwide relenting phases take several days to complete and a regular observable anomaly occurs in the price distribution in the midst of every one. Over the course of several days, there emerges two distinct sets of stations with prices on average 7 to 11.6 cpl apart from one another (25 to 42 U.S. cents per gallon). The set of high priced stations grows and the set of low priced stations shrinks until the latter is empty and the relenting phase is complete.¹¹ The Spike and Buy decision rule calls for purchasing from a station in the lower priced set once the bimodal distribution is observed. The rule does not depend on any critical values, and thus the Spike and Buy rule is the only member of its class.

Myopic Decision Rule

Finally, we calculate prices under a myopic decision rule. This is not a strategic rule – it is essentially purchasing randomly without regard to price – and is a benchmark against which other strategic classes can be compared.

¹¹It is well known that at an individual station, the full trough-to-peak price increase happens in a single shot. See, for example, Noel (2007b) or Wang (2009).

4.3 A Preliminary Look at Trough Predictability

Decision rules work essentially by predicting when the cycle troughs will occur. The penalty for missing the trough can be significant. Amplitudes average 7 to 11.6 cpl (the cycle amplitude) on average across cities and are often as high as 18 cpl (65 U.S. cents per gallon). Purchasing at the troughs instead of myopically typically reduces prices by close to half the cycle amplitude.¹²

We examine the potential for trough predictability using the probit regression of Equation 6, but replacing $(I_{mt} = R \mid I_{m,t-1} = U)$ with $I_{m,t+1} = R \mid I_{mt} = U$ as the dependent variable (since we want to predict the trough just before prices start to rise). On the right hand side, in the W_{mt} terms, we include the predictor variables, τ_{mt} , that form the bases for the predictions in each class. The set of predictors are $\tau_{mt} = \{POSITION_{mt}, DECREASE_{mt}, DAYS_{mt}, DOW_t^2 \dots DOW_t^7, \}$, where $DOW_t^1 = SUNDAY$ is the omitted day of the week indicator. The four predictors (taking the DOW_t^d dummies as one) correspond to the trigger mechanisms associated with the first four decision rule classes above. A significant coefficient on a predictor means that the associated decision rule class has potential power in predicting troughs. (The critical values μ_c can be subsequently fine tuned.)

The Spike and Buy rule is different in that it is based on the idea that price increases by the first handful of stations will likely be followed by more widespread price increases the following day, thus identifying the trough. One way to get at this is to use Equation 6, replacing $(I_{mt} = R \mid I_{m,t-1} = U)$ with $(I_{m,t+1} = R \mid I_{mt} = R)$ and putting $\tau_{mt} = \{I_{m,t-1} = U\}$ into the W_{mt} . In other words, we test if it is more or less likely for prices to rise in period $t + 1$, when they rise for the first time in period t . If so, the first price increase begets more price increases and the first price increases at a handful of stations indicate more price increases by the majority are to come.

<Table 3 placed about here>

Table 3 presents results from the preliminary probit analysis. It shows that new relenting phases $(I_{m,t+1} = R \mid I_{m,t} = U)$ are more likely to occur when *POSITION* is lower (i.e. margins are lower), when *DECREASE* is lower (i.e. when prices have fallen by a larger amount in absolute value since the last peak), and when *DAYS_{mt}* is greater (i.e. when more time has elapsed since the

¹²The requirement that agents must purchase at least every w days limits these gains.

last peak). The results are uniformly significant across all cities. The table also shows that troughs over the whole sample are most concentrated Tuesdays to Thursdays, Wednesdays especially, and for a few cities Fridays as well. Weekends and Mondays are least likely. The final row shows that, conditional on a price increase in period t , a price increase is more likely in $t + 1$ if there was a price decrease in period $t - 1$ (i.e. a second price increase is more likely than a decrease). Simply put, relenting phases take several days to complete, so the first stations' price increases indicate more widespread price increases to come, and indicate the still low priced stations are at the trough.¹³

4.4 Effectiveness and Stability of Decision Rule Classes

The results of the preliminary analysis suggest that all five classes of decision rules are potentially valuable in predicting troughs, and we carry them forward. We now calculate prices actually obtainable under various decision rules and associated critical values μ_c , and compare them to myopic purchasing. We do not yet address how agents would know to use these rules, we only wish to establish which rules and classes would be beneficial if used.

We simulate purchase profiles for agents under each decision rule within each class and over each time interval. We use half-year intervals throughout our empirical analysis, though in the subsequent figures we show annual intervals for readability. We compare the prices paid under each decision rule with the prices paid under a myopic purchasing rule, and calculate the price reduction as the difference.

Figures 3 through 6 plot reduction profiles for the city of Melbourne - price reductions plotted against the sequence of μ_c , one for each class of decision rules. The heavy black curve in each figure shows price reductions for the chosen μ_c for the whole sample period, if μ_c were used throughout. Other curves plot price reductions against different μ_c on an annual basis. Confidence intervals are suppressed. Figures for other cities are excluded to conserve space but are qualitatively similar.

Figure 3 shows the Position Based class. The optimal value of μ_{PB} , the critical position value, is 3 cpl over the whole sample and would have resulted in a statistically significant price reduction of 2.4 cpl. Values of μ_{PB} between 2.5 and 5.5 cpl yield only slightly lower reductions overall. In

¹³Consistent with this, we find the second (market average) price increase in period $t + 1$ is larger than the first increase in period t . In other words, the first set of stations to increase prices trough-to-peak is a relatively small subset and most hold off to the second day, creating an easy purchasing window.

individual years, however, optimal values of μ_{PB} fluctuate, ranging from -1.5 cpl in 2013 to 5.5 cpl in 2010. Even though the optimum is $\mu_{PB} = 3$ overall, the annual reduction profiles show that that μ_{PB} would not result in any price reductions in the last few years and there is no reason to suspect that value would be useful going forward.

Figure 4 shows the Spike and Drop Class. The optimal value of μ_{SD} , the cumulative price decrease since the last peak, is 8 cpl over the whole sample. It results in a statistically significant price reduction of 2.4 cpl. Values of μ_{SD} between 6.5 and 9 cpl yielded only slightly lower reductions overall. Optimal values of μ_{SD} rose in individual years from 7 cpl in 2007 to 16.5 cpl in 2013. If one were to have used the overall optimum of $\mu_{SD} = 8$ in 2013, it would have result in a price *increase* of 2 cpl relative to myopic purchasing. That is, paying no attention to prices in 2013 would have been better than using Spike and Drop's optimum over the sample.

Figure 5 shows the Spike and Wait Class. The best value of μ_{SD} , the number of days since the last peak, is 5 or 12 over the whole sample and results in a statistically significant price reduction of 2 cpl. A value of $\mu_{SD} = 4$ or 11 also yields price reductions, but a third lower. Intermediate values yield either no price reductions or price increases relative to myopic purchasing. The virtual duality of optimality stems from the fact that cycle periods were often a week.¹⁴ Using the overall optimum of $\mu_{SD} = 5$ would have resulted in price *increases* in each of the last few years, relative to myopic purchasing.

Finally, Figure 6 shows price reductions for the Calendar Based Class. The optimal value of μ_{CB} is Wednesday and would result in a statistically significant price reduction of 2 cpl. However, cycle troughs rarely occur on Wednesdays towards the end of the sample. If one were to rely on the overall optimal value of μ_{CB} of Wednesday, it would have resulted in no significant price reductions in the final three years, relative to myopic purchasing. In fact, the overall price reduction of 2 cpl is an average of 4 cpl in earlier years and 0 cpl in later years. The decision rule has lost its effectiveness over time.

We do not display a figure for the fifth class of decision rules, the Spike and Buy Class, as it is a singleton class. Price reductions ranged from 3.7 cpl to 5.2 cpl per year, with an average on 4.9 cpl (17.7 U.S. cents per gallon).

¹⁴For smaller values of w , the shorter 5 days strictly dominates 12 days.

The results taken together show that optimal decision rules change over time so regular updating of forecasts is important. This can also be seen in the variation in the annual curves in Figures 3 to 6. What was optimal last year may not be optimal this year.

We define class stability as the ability to reliably yield price reductions over time with periodic reoptimization. A class is less stable if the optimal rule within the class tends to "bounce around" a lot, meaning that past optimizations hold little information about future optimal values. It is also less stable if deviations from optimality lead to much inferior outcomes, shown as greater concavity in the reductions profiles. As different classes are based on different characteristics, stability can vary across classes. The more stable a class, the more effective it is with regular updating of forecasts.

Table 4 gives an example of how stability affects effectiveness. We show the price reductions, relative to myopic purchasing, that would be obtained over the last half of the sample, 2010H2 (second half of 2010) to 2013H1 (first half of 2013), under two extreme assumptions. The first is the assumption of no updating. For this, we assume that the optimal values calculated in the first six months of the sample (2007H1) are learned and then used over time without further updating.¹⁵ The second assumption is that of perfect foresight (bottom panel). The difference between the two yields a measure of how much decision rules lost their effectiveness over time relative to the gains that were potentially available.

The top panel of the table shows that for each city and four out of five classes, the price reductions tail off significantly as the rules become stale. Vertical classes - the Position Based and Spike and Drop decision classes - lose between one-third and all of their effectiveness by the last half of the sample (top panel), relative to the reductions under perfect foresight (bottom panel). Horizontal rules - Spike and Wait and Calendar Based rules - fare worse, losing between one-half and all of their effectiveness. For most classes and cities, there are no significant reductions at all compared to myopic purchasing. In a word, the rules became useless. The Noel (2012) assumption eventually loses strength. The one exception is the Spike and Buy decision rule which, as the only member of its class, maintains its stability by construction. It is not surprising in light of the fact

¹⁵More accurately, this is "no updating" after the first update (from 2007H1). This is subtly different than no forecasting - no forecasting is equivalent to myopic purchasing, i.e. the agent has no information on what to do and never learns the optimal decision rules in any period.

it does not depend on cycle amplitudes or periods but rather on a guaranteed regular anomaly in the price distribution during relenting phases. It also yields the highest price reductions of all in both the no updating (top panel) and perfect foresight (bottom panel) cases.

4.5 Forecasting Algorithms

The "no updating" algorithm is not an effective one for forecasting price reductions for most classes. Conversely, perfect foresight is highly effective but not at all realistic.

What happens if we rule out perfect foresight? Can effective future decision rules be meaningfully forecast? Which forecasting algorithms are best? How close can one get to replicating the perfect foresight outcome? Or can significant price reductions be obtained at all? In this section, we consider a series of feasible forecasting algorithms based on first order and higher order Markov chains with a goal of forecasting troughs only on past and known information.

Let Φ^c denote the vector of available decision rules in class c , and a member of that class be ϕ . A first order Markov forecasting process takes the form, $\Phi_q^c = \Gamma\Phi_{q-1}^c$ for period q where Γ is a transition matrix. Suppressing the superscripts, the no updating (after the first interval) case, discussed above, is the chain where Γ is the identity matrix, i.e. do this period the same as last period.

We consider a full adjustment, memoryless, forecasting algorithm. This is among the simplest. Let ϕ_{q-1}^* be the ex-post optimal decision rule within the class in period $q - 1$, that is:

$$\phi_{q-1}^* = \arg \min_{\phi_{q-1}} \left(\sum_{t \in (q-1)} \bar{p}_{q-1} \mid \phi_{q-1} \right), \phi_{q-1} \in \Phi_{q-1}, \quad (11)$$

where \bar{p} is the mean price. Regardless of whether or not ϕ_{q-1}^* is actually used in $q - 1$, the agent simply implements the choice $\phi'_q = \phi_{q-1}^*$ in period q irrespective of his own past behavior. In other words, the agent learns which decision rule was optimal last time and uses it this time. The matrix Γ is simply a row of ones in the row corresponding to decision rule ϕ^* and zero elsewhere.

We also consider a series of partial adjustment forecasting algorithms. In these cases, the agent implements ϕ'_q where $\phi'_q = \kappa\phi_{q-1}^* + (1 - \kappa)\phi'_{q-1}$ for $\kappa \in [0, 1]$. The weighted average is taken by averaging the relevant values of μ_c for class c . For example, if the agent chose a critical position

value of $\mu_{PB} = 3.5$ in period $q - 1$ and ex post it is learned that $\mu_{PB} = 2.5$ was optimal in $q - 1$, the agent chooses $\mu_{PB} = 3$ in period q . Partial adjustment can protect agents against any large swings back and forth in the optimal critical values, which would repeatedly push the agent down the slope of the reduction profile away from the optimum. We report results for $\kappa = 0.5$, with results for other κ 's as expected. Note that $\kappa = 1$ corresponds to full adjustment and $\kappa = 0$ corresponds to no updating (i.e. adopting the optimal decision rules from the first half-year of the sample and then no updating thereafter).

We also experiment with various trend extrapolation algorithms. Essentially, we take $\phi'_q = \phi_{q-1}^* + \varkappa(\phi_{q-1}^* - \phi_{q-2}^*)$. A linear trend corresponds to $\varkappa = 1$. For example, if optimal critical values were $\mu_{PB} = 4.5$ and $\mu_{PB} = 4$ in periods $q - 2$ and $q - 1$ respectively, the agent simply chooses $\mu_{PB} = 3.5$ in period q under a linear trend. This allows agents to identify trends in maxima and attempt to extrapolate forward.

Tables 5A and 5B report our results on price reductions relative to myopic purchasing for four forecasting algorithms. The top panel of Table 5A shows the perfect foresight algorithm. While not feasible, it is the best case scenario and the benchmark case to which we compare feasible algorithms.¹⁶ The bottom panel of Table 5A shows the full adjustment algorithm, the top panel of 5B shows the partial adjustment algorithm with $\kappa = 0.5$, and the bottom panel of 5B shows the linear trend extrapolation algorithm with $\varkappa = 1$.

We begin with the full adjustment forecasting algorithm. This is the most straightforward feasible algorithm – basically, if it was the best last time, do it this time. We highlight several key results.

First, we find that even in the absence of perfect foresight, and even with significant changes in the cycle over time, cycle troughs are still readily predictable. We find significant price reductions relative to myopic purchasing (i.e. no forecasting) for almost every class-city combination (Table 5A, bottom panel).

For the Position Based Class, the full adjustment algorithm yields price reductions between 0.8 and 3.1 cpl across markets relative to myopic purchasing. It is significant in four of five markets.

¹⁶The price reductions for perfect foresight given here differ from those reported in Table 4 in that they are calculated over the entire sample except the first year (which is used to form initial forecasts), and not only the last three years.

The Spike and Drop Class yields significant price reductions between 1.1 and 2.5 cpl. The Spike and Wait Class yields significant price reductions of 2.9 to 3.3 cpl and the Calendar Based Class yields significant price reductions of 1.6 to 2.5 cpl across markets. Of these classes, the Spike and Wait Class is the most effective, yielding an average price reduction equivalent to 11.1 US cents per gallon overall. The most effective class of all is the Spike and Buy Class, which yields price reductions of 3.3 cpl to 5.4 cpl across markets, equivalent to 15.3 US cents per gallon on average.

Second, although we find significant price reductions, they are lower than those assumed under perfect foresight, for four of five classes. In other words, as cycles evolve, optimal decision rules change and old methods do not perfectly replicate old results. Only the Spike and Buy Class can fully replicate price reductions under perfect foresight, but this is by virtue of being a singleton class.

It is interesting to compare how much of the "available" price reductions with perfect foresight (Table 5A top panel) are still captured with the feasible full adjustment algorithm (bottom panel) for the other four classes. We find the Spike and Wait Class holds up well using the full adjustment algorithm. Price reductions are about 80% relative to what they would be with perfect foresight. Cycle periods tend to change gradually and agents do not slide too far down the reduction profiles from one interval to the next. The Calendar-Based Class, the other horizontal class, retains 86% relative to perfect foresight, but is less effective overall. In contrast, the vertical classes - the Position Based Class and the Spike and Drop Class - retain only 59% and 68% of their value on average across cities, respectively. They also yield lower price reductions than Spike and Wait. Cycle amplitudes tend to be more erratic in the short term, leading to difficulty in forecasting troughs based on them.

We have calculated price reductions assuming the agent chooses and exclusively remains with a particular class. But what if agents were allowed to switch back and forth between classes over time? There is no reason in principle why agents must stick to just one class. One might expect the extra flexibility would improve upon our estimates, but surprisingly, this is not true.

The last row of Table 5A reports price reductions under this kind of flexibility. Not surprisingly, we find that agents with full flexibility would maximize price reductions by using Spike and Buy in every interval, so we do not repeat these same results in the table. Instead, we ask what would

happen if agents were flexible among the remaining four classes. In three cities, the price reductions are a tiny amount higher (a tenth of a cent at most) compared to using the Spike and Wait Class exclusively. In two cities, they are actually lower.

How can more flexibility be harmful? If agents had perfect foresight, it cannot. But with feasible forecasting, it can. When an agent is coaxed into switching to a relatively unstable vertical rule because it turned out to be optimal in a previous interval, only to find it reverts back to performing poorly again in the next, the agent slides down the (concave) reduction profile and outcomes suffer. As a result, constraining oneself to the more stable Spike and Wait Class is not only simpler but at least as good if not better than exercising flexibility across classes.

We now turn to other feasible algorithms. The top panel of Table 5B reports price reductions from the partial adjustment forecasting algorithm. In spite of the additional complexity, the results for the Position Based, Spike and Drop and Calendar Based Classes are all very similar to the full adjustment case, usually within two-tenths of a cent. The largest improvement is half a cent (Position Based Class in Sydney) and the largest deterioration is three-tenths of a cent (Spike and Drop Class in Brisbane). The results for the Spike and Wait Class, however, fall dramatically. The reason is that there are large increases in the period of the cycle, which under partial adjustment results in followers missing the troughs and making purchases near the peak for several years. The partial adjustment model takes not fully account for known information about the new periodicity. Of all classes, the Spike and Buy Class yields the largest price reductions.

The bottom panel of Table 5B presents results under the linear trend forecasting algorithm. In a word, the results are poor. The price reductions for the four vertical and horizontal classes are uniformly lower than those under full adjustment. We experiment with various types of more complicated trend forecasting with similarly poor results. Optimal values can fluctuate back and forth from interval to interval and typically the effect of linear trend forecasting is to push agents too often into the tails of the reduction profiles.

We conclude that the most effective feasible algorithm is the full adaptation model. We also conclude that the most effective classes for use under that model are the Spike and Buy and Spike and Wait Classes.

4.6 Discussion

In terms of implementation, there are pros and cons of both the Spike and Buy Class and the Spike and Wait Class and different agents are likely to prefer different ones. The Spike and Buy Class yields greater price reductions but requires more flexibility than other classes since agents do not typically know when the purchase window will open until usually the same day and then they need to act at once. The Spike and Wait Class in contrast can be planned more easily as the day for purchase is determined many days in advance, e.g. 5 to 12 days after the last peak. Given agents are heterogeneous in their schedule flexibility, a mix of usage of these classes would likely result.

It is clear that consumers are interested in timing the troughs. A 2007 survey commissioned by the ACCC (Australian Competition and Consumer Commission) of 775 people in the five major Australian cities found consumers were largely aware of the cycle – 83% of consumers in the four eastern cities were aware of the cycles and 61% were aware of the somewhat less regular cycles in Perth. Cycles were then synchronized to the day of the week in the eastern cities. The survey found that a majority of consumers attempted to regularly time purchases to be near the troughs – 59% in those cities (ACCC (2007)).

Consumers not only showed an interest in timing the troughs but were also effective in actually doing so. Examining the 2009-2010 period during which cycles in eastern cities were synchronized to the day of the week, ACCC (2010) reports that more than twice as many consumers purchased at the trough (on Wednesdays) than at or near the peak (on Fridays). The data was collected directly from firms. In both Sydney and Melbourne, 24% of volumes came on trough day, compared to just 11% on peak day. In Adelaide, it was 27% versus 10%, and in Brisbane it was 23% versus 12%. In Perth, where there were still day-of-the-week synchronized cycles as of 2013, and ACCC (2013) reports that 23% of volumes came on the Wednesday trough but only 15% on the Thursday peak. Volumes in the few days before the trough were also relatively high.

So both the interest and actual execution of moving purchases around is evident. The fact that cycles were synchronized to the calendar in the past meant that Calendar-Based decision rules were particularly simple and consumers used them effectively. This article makes the important point that, even in the absence of day-of-the-week synchronicity, cycle troughs are still readily

predictable. There are still significant long term gains to be had with the Spike and Buy and Spike and Wait classes of decision rules.

Because these rules are not quite as simple of day-of-the-week purchasing, there is a potentially greater role for government and private industry in helping consumers find the troughs. Either sector would be well positioned to perform the forecasting function on behalf of consumers, and give trough predictions and recommended purchase windows to interested consumers. The ACCC maintains a popular price information website which lists prices in each city, with graphs, and for years reported the high and low priced days of the calendar week. Media outlets widely report on the price cycles and refer consumers back to the ACCC website. Also, the popular private website Motormouth, the Australian equivalent of GasBuddy in North America, provides consumers with weekly and monthly graphs of the cycles (in addition to local station prices) that clearly show the cycles. These agencies could perform the forecasting function and suggest purchase windows under the Spike and Buy or Spike and Wait classes, both using their websites and more directly through mobile apps, email alerts, social media feeds, etc., alerting consumers of a good time to buy.

While the widespread adoption of such an app or alert system could in principle affect the cycles, it is far from clear that it would destabilize them. After all, the cycles continued to perform strongly even when 59% of consumers were both aware of them and trying to time them, and even when volumes at the troughs were routinely double volumes at the peaks.

Finally, it is worth noting the connection between the Calendar-Based class and the Spike and Wait class. When cycle troughs line up with days of the week, the two classes are perfectly correlated. When this occurs, it can be simpler to convey to consumers expected trough days in terms of days of the week. However, the results of this article show that, even when there are deviations of the cycle period from fixed days on the calendar, the Spike and Wait Class continues to be robust and is therefore a superior metric. Interestingly, the ACCC continued to report "low price days of the week" on its website for years after the Calendar Based Class became ineffective as a predictive tool, while the Spike and Wait class would have consistently continued to find gains.

5 Conclusion

We conclude that, of all the feasible forecasting algorithms explored, the best is also the simplest. The full adjustment forecasting algorithm produces results closest to perfect foresight and yields the highest price reductions relative to myopic purchasing. The best guess of what will be optimal next time is just what is optimal this time. More complicated algorithms studied do not improve upon this result.

We conclude that the greatest price reductions under full adjustment occur with the Spike and Buy and Spike and Wait Classes of decision rules. Price reductions under Spike and Buy and under Spike and Wait are approximately 15.3 and 11.1 U.S. cents per gallon, respectively, on average across the five markets. The estimates derive from forecasts that use only prior and known information and do not assume perfect foresight, as with previous work.

An interesting feature of our application is that the loss function is significantly asymmetric so that misforecasting the trough a day too early has little consequences, but misforecasting it a day too late is severe. This in turn affects how long agents are willing to "push their luck", which in turn limits potential price reductions. However, we find that cycle troughs are still predictable and price reductions are significant. In other words, price reductions are available not just in principle but also in practice.

This article adds to our understanding of how to forecast gasoline prices in the short run and, more generally, how to forecast outcomes from asymmetric processes with *feasible* forecasting algorithms. The results also have a potentially useful policy application. It can be used by government or private agencies interested in learning how to forecast cycle troughs for the purpose of sharing that information with consumers. Many jurisdictions have shown an interest in this. In Western Australia (where Perth is located) and also in Austria, governments have legislated mandatory price pre-notification laws that require firms to announce price increases a day in advance (Dewenter and Heimeshoff (2012) and Haucap and Muller (2012)). A similar law was considered on a nationwide basis in Australia and Canada (Noel (2012)). This article offers an alternative. Forecasting in the manner outlined here can be undertaken by a central agency or by private firms and the resulting purchase windows updated and relayed to consumers on an ongoing basis. While volatility in gaso-

line prices will surely remain a controversial policy topic for some time to come, the lessons of our article can help in bringing some additional transparency to this process.

6 References

Atkinson, B. (2009). “Retail Gasoline Price Cycles: Evidence from Guelph, Ontario Using Bi-Hourly, Station-Specific Retail Price Data”, *Energy Journal* 30, pp. 85-109.

Australian Competition and Consumer Commission (2007). “Petrol Prices and Australian Consumers: Report of the ACCC Inquiry into the Price of Unleaded Petrol”, Canberra.

Australian Competition and Consumer Commission (2010). “Monitoring of the Australian Petroleum Industry”, Canberra.

Australian Competition and Consumer Commission (2013). “Monitoring of the Australian Petroleum Industry”, Canberra.

Castanias, R. and H. Johnson (1993). “Gas Wars: Retail Gasoline Price Fluctuations”, *Review of Economics and Statistics* 75, 171-174.

De Roos, N. and H. Katayama (2013). “Gasoline Price Cycles Under Discrete Time Pricing”, *Economic Record* 89, pp. 175-193.

Dewenter, R. and U. Heimeshoff (2012). “Less Pain at the Pump? The Effects of Regulatory Interventions in Retail Gasoline Markets”, Dusseldorf Institute for Competition Economics (DICE) Working Paper 51.

Doyle, J., E. Muehlegger and K. Samphantharak (2010). “Edgeworth Cycles Revisited”, *Energy Economics* 32, pp. 651-660.

Eckert, A. (2002). “Retail Price Cycles and Response Asymmetry”, *Canadian Journal of Economics* 35, pp. 52-77.

Eckert, A. (2003). “Retail Price Cycles and the Presence of Small Firms”, *International Journal of Industrial Organization* 21, pp. 151-170.

Edgeworth (1925). “The Pure Theory of Monopoly”, in *Papers Relating to Political Economy*, Vol. 1. London: MacMillan, pp. 111-142.

Gicheva, D., J. Hastings, and S. Villas-Boas (2010). “Investigating Income Effects in Scanner Data: Do Gasoline Prices Affect Grocery Purchases?”, *American Economic Review Papers and Proceedings* 100, pp. 480-484.

Haucap, J. and H.C. Muller (2012). “The Effects of Gasoline Price Regulations: Experimental

Evidence”, Dusseldorf Institute for Competition Economics (DICE) Working Paper 147.

Lewis, M. (2009). “Temporary Wholesale Gasoline Price Spikes have Long Lasting Retail Effects: The Aftermath of Hurricane Rita”, *Journal of Law and Economics* 52(3), pp. 581-606.

Lewis, M. (2012). “Price Leadership and Coordination in Retail Gasoline Markets with Price Cycles”, *International Journal of Industrial Organization* 30, pp. 342-351.

Lewis, M. and M. Noel. (2011). “The Speed of Gasoline Price Response in Markets with and without Edgeworth Cycles”, *Review of Economics and Statistics* 93:2, pp. 672-682.

Hughes, J.E., C.R. Knittel, and D. Sperling (2008). “Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand”, *Energy Journal* 29, pp. 113-134.

Maskin, E. and J. Tirole (1988). “A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles”, *Econometrica* 56, pp. 571-599.

Noel, M.D. (2002). “Edgeworth Price Cycles in Retail Gasoline Markets”, MIT Dissertation.

Noel, M.D. (2007a). “Edgeworth Price Cycles, Cost-Based Pricing, and Sticky Pricing in Retail Gasoline Markets”, *Review of Economics and Statistics* 89, pp. 324-334.

Noel, M. D. (2007b). “Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market”, *Journal of Industrial Economics* 55, pp. 69-92.

Noel, M.D. (2008). “Edgeworth Price Cycles and Focal Prices: Computational Dynamic Markov Equilibria”, *Journal of Economics and Management Strategy* 17, pp. 345-377.

Noel, M.D. (2009). “Do Retail Gasoline Prices Respond Asymmetrically to Cost Shocks? The Influence of Edgeworth Cycles”, *RAND Journal of Economics* 40, pp. 582-595.

Noel, M.D. (2012). “Edgeworth Price Cycles and Intertemporal Price Discrimination”, *Energy Economics* 34, pp. 942-954.

Noel, M.D. (2015). “Do Edgeworth Price Cycles Cause Higher or Lower Prices?”, working paper.

Wang, Z. (2009a). “Station Level Gasoline Demand in an Australian Market with Regular Price Cycles”, *Australian Journal of Agricultural and Resource Economics* 53, pp. 467-483.

Wang, Z. (2009b). “(Mixed) Strategy in Oligopoly Pricing: Evidence from Gasoline Price Cycles Before and Under a Timing Regulation”, *Journal of Political Economy* 117, pp. 987-1030.

Zimmerman, P.R., J.M.Yun and C. T. Taylor (2013). “Edgeworth Price Cycles in Gasoline:

Evidence from the U.S.”, *Review of Industrial Organization* 42, pp. 297-320.

Table 1. Within-Phase Price Changes and Phase Transition Probabilities

	<u>Adelaide</u>	<u>Brisbane</u>	<u>Melbourne</u>	<u>Perth</u>	<u>Sydney</u>
<u>Within-Phase Price Changes</u>					
$\alpha^R = E(\Delta \text{RETAIL}_{mt} I_{mt} = "R")$ (average price change - relenting)	4.864** (0.137)	3.090** (0.082)	3.819** (0.101)	2.921** (0.089)	3.792** (0.105)
$\alpha^U = E(\Delta \text{RETAIL}_{mt} I_{mt} = "U")$ (average price change - undercutting)	-1.835** (0.028)	-1.364** (0.017)	-1.564** (0.019)	-1.192** (0.015)	-1.556** (0.022)
<u>Phase Transition Probabilities</u>					
λ^{RR} (R to R)	0.583** (0.019)	0.624** (0.018)	0.606** (0.019)	0.584** (0.020)	0.604** (0.019)
λ^{RU} (R to U)	0.417** (0.019)	0.376** (0.018)	0.394** (0.019)	0.416** (0.020)	0.396** (0.019)
λ^{UR} (U to R)	0.159** (0.009)	0.168** (0.009)	0.163** (0.009)	0.175** (0.010)	0.164** (0.009)
λ^{UU} (U to U)	0.841** (0.009)	0.832** (0.009)	0.837** (0.009)	0.825** (0.010)	0.836** (0.009)

Standard errors in parentheses. *Significant at the 5% level. **Significant at the 1% level.

Table 2. Cycle Characteristics

	<u>Adelaide</u>	<u>Brisbane</u>	<u>Melbourne</u>	<u>Perth</u>	<u>Sydney</u>
Relenting Phase Duration	2.396** (0.111)	2.661** (0.127)	2.540** (0.120)	2.403** (0.117)	2.524** (0.119)
Undercutting Phase Duration	6.286** (0.349)	5.942** (0.327)	6.140** (0.341)	5.702** (0.329)	6.095** (0.337)
Cycle Period	8.681** (0.366)	8.602** (0.351)	8.680** (0.361)	8.105** (0.349)	8.619** (0.357)
Cycle Amplitude	11.652** (0.651)	8.220** (0.465)	9.702** (0.549)	7.019** (0.391)	9.570** (0.532)
Cycle Asymmetry	2.624** (0.190)	2.233** (0.163)	2.417** (0.176)	2.372** (0.179)	2.415** (0.176)

Durations and cycle periods in days, amplitudes in cents per liter, asymmetries are unit free.
Standard errors in parentheses. *Significant at the 5% level. **Significant at the 1% level.

Table 3. Potential Value of Decision Rule Classes in Predicting Troughs

	<u>Adelaide</u>	<u>Brisbane</u>	<u>Melbourne</u>	<u>Perth</u>	<u>Sydney</u>
Dependent Variable = ($I_{m,t+1} = R \mid I_{mt} = U$)					
POSITION	-0.030** (0.002)	-0.033** (0.002)	-0.034** (0.002)	-0.049** (0.003)	-0.034** (0.002)
DECREASE	-0.023** (0.002)	-0.034** (0.003)	-0.026** (0.002)	-0.076** (0.005)	-0.035** (0.002)
DAYS	0.023** (0.003)	0.031** (0.004)	0.011** (0.002)	0.040** (0.005)	0.024** (0.003)
MONDAY	0.043 (0.022)	0.030 (0.017)	0.006 (0.014)	0.008 (0.017)	0.002 (0.018)
TUESDAY	0.245** (0.030)	0.237** (0.027)	0.224** (0.026)	0.263** (0.031)	0.238** (0.029)
WEDNESDAY	0.312** (0.036)	0.437** (0.034)	0.390** (0.033)	0.642** (0.036)	0.397** (0.036)
THURSDAY	0.184** (0.038)	0.107** (0.031)	0.160** (0.033)	0.303** (0.061)	0.098** (0.033)
FRIDAY	0.031 (0.023)	0.130** (0.030)	0.157** (0.033)	-0.028* (0.014)	0.101** (0.029)
SATURDAY	-0.014 (0.018)	0.060** (0.021)	0.062** (0.020)	-0.032** (0.012)	0.010 (0.019)
Dependent Variable = ($I_{m,t+1} = R \mid I_{mt} = R$)					
$I_{m,t-1} = U$	0.517** (0.032)	0.580** (0.024)	0.626** (0.025)	0.507** (0.033)	0.567** (0.028)

Standard errors in parentheses. *Significant at the 5% level. **Significant at the 1% level.

Table 4. Stability of Decision Rule Classes Without Updating

	<u>Adelaide</u>	<u>Brisbane</u>	<u>Melbourne</u>	<u>Perth</u>	<u>Sydney</u>
<u>Price Reductions 2010H2-2013H1 Based on 2007H1 Optimal Decision Rules</u>					
Position Based	2.464** (0.604)	1.166* (0.532)	1.106 (0.622)	-0.084 (0.509)	1.121 (0.596)
Spike and Drop	1.995** (0.475)	1.416** (0.486)	-0.546 (0.506)	0.135 (0.517)	0.112 (0.662)
Spike and Wait	1.710** (0.504)	0.364 (0.543)	0.278 (0.537)	2.210** (0.363)	0.530 (0.529)
Calendar Based	0.228 (0.514)	-0.106 (0.468)	-0.072 (0.537)	-0.428 (0.358)	0.110 (0.477)
Spike and Buy	5.781** (0.443)	3.901** (0.427)	4.358** (0.486)	3.768** (0.364)	4.063** (0.408)
<u>Price Reductions 2010H2-2013H1 with Perfect Foresight</u>					
Position Based	3.664** (0.549)	1.960** (0.483)	2.899** (0.539)	2.586** (0.386)	2.818** (0.478)
Spike and Drop	3.595** (0.606)	2.215** (0.470)	2.811** (0.541)	3.135** (0.382)	2.323** (0.523)
Spike and Wait	3.352** (0.511)	2.363** (0.442)	2.912** (0.491)	3.421** (0.366)	2.903** (0.435)
Calendar Based	1.268* (0.541)	0.680 (0.442)	0.801 (0.523)	3.488** (0.360)	0.856 (0.459)
Spike and Buy	5.781** (0.443)	3.901** (0.427)	4.358** (0.486)	3.768** (0.364)	4.063** (0.408)

Standard errors in parentheses. *Significant at the 5% level. **Significant at the 1% level.

Table 5A. Price Reductions Under Forecasting Algorithms

	<u>Adelaide</u>	<u>Brisbane</u>	<u>Melbourne</u>	<u>Perth</u>	<u>Sydney</u>
<u>Perfect Foresight</u>					
Position Based	4.048** (0.649)	2.287** (0.509)	3.057** (0.543)	2.194** (0.292)	3.069** (0.531)
Spike and Drop	3.568** (0.669)	2.526** (0.550)	3.220** (0.579)	2.417** (0.281)	2.931** (0.560)
Spike and Wait	3.919** (0.536)	2.753** (0.471)	3.469** (0.506)	3.013** (0.257)	3.529** (0.484)
Calendar Based	2.732** (0.550)	1.984** (0.477)	2.260** (0.517)	2.907** (0.260)	2.443** (0.492)
Spike and Buy	5.420** (0.525)	3.778** (0.479)	4.402** (0.506)	3.308** (0.257)	4.346** (0.477)
<u>Full Adjustment Algorithm</u>					
Position Based	3.139** (0.600)	0.846 (0.520)	1.873** (0.571)	1.655** (0.312)	1.464** (0.533)
Spike and Drop	2.514** (0.572)	1.089* (0.523)	2.414** (0.524)	2.104** (0.300)	1.988** (0.510)
Spike and Wait	3.335** (0.542)	2.286** (0.461)	2.951** (0.511)	2.876** (0.261)	3.181** (0.484)
Calendar Based	2.506** (0.543)	1.666** (0.487)	1.859** (0.524)	2.298** (0.278)	2.179** (0.493)
Spike and Buy	5.420** (0.525)	3.778** (0.479)	4.402** (0.506)	3.308** (0.257)	4.346** (0.477)
All Classes (except Spike & Buy)	3.372** (0.579)	2.411** (0.471)	2.622** (0.515)	2.791** (0.257)	3.194** (0.481)

Standard errors in parentheses. *Significant at the 5% level. **Significant at the 1% level.

Table 5B. Price Reductions Under Forecasting Algorithms

	<u>Adelaide</u>	<u>Brisbane</u>	<u>Melbourne</u>	<u>Perth</u>	<u>Sydney</u>
<u>Partial Adjustment Algorithm</u>					
Position Based	3.156** (0.587)	0.697 (0.505)	1.872** (0.528)	1.471** (0.323)	1.906** (0.525)
Spike and Drop	2.738** (0.565)	0.704 (0.519)	2.403** (0.498)	1.801** (0.339)	1.850** (0.541)
Spike and Wait	-0.181 (0.565)	0.696 (0.468)	0.183 (0.521)	-0.509* (0.257)	-0.287 (0.505)
Calendar Based	2.434** (0.538)	1.617** (0.489)	1.865** (0.523)	2.511** (0.282)	2.117** (0.492)
Spike and Buy	5.420** (0.525)	3.778** (0.479)	4.402** (0.506)	3.308** (0.257)	4.346** (0.477)
<u>Trend Forecasting Algorithm</u>					
Position Based	2.088** (0.596)	0.313 (0.552)	1.043 (0.628)	1.131** (0.299)	0.509 (0.533)
Spike and Drop	1.817** (0.601)	0.292 (0.553)	1.785** (0.551)	0.772* (0.328)	1.142* (0.532)
Spike and Wait	1.696** (0.548)	1.255** (0.464)	2.166** (0.518)	1.192** (0.261)	2.633** (0.485)
Calendar Based	1.454** (0.549)	1.377** (0.483)	1.496** (0.519)	0.920** (0.267)	1.251* (0.492)
Spike and Buy	5.420** (0.525)	3.778** (0.479)	4.402** (0.506)	3.308** (0.257)	4.346** (0.477)

Standard errors in parentheses. *Significant at the 5% level. **Significant at the 1% level.

Figure 1. Retail and Rack Prices, Melbourne, Jul.-Dec. 2010

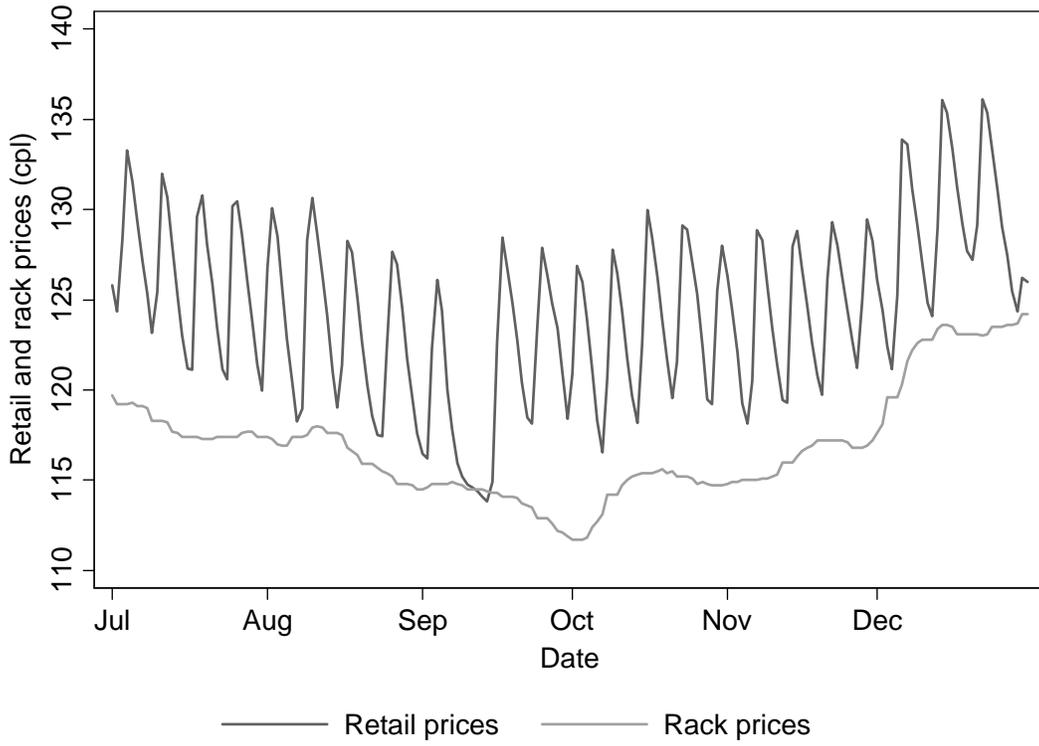


Figure 2. Theoretical Example of an Edgeworth Price Cycle

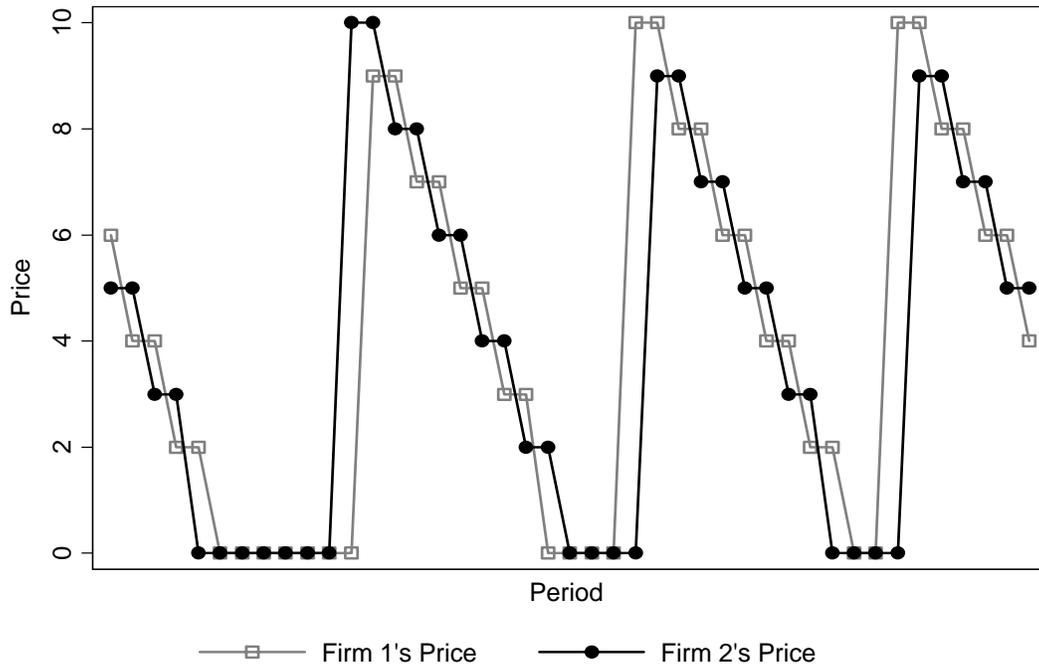


Figure 3. Reduction Profile for the Position-Based Class, Melbourne

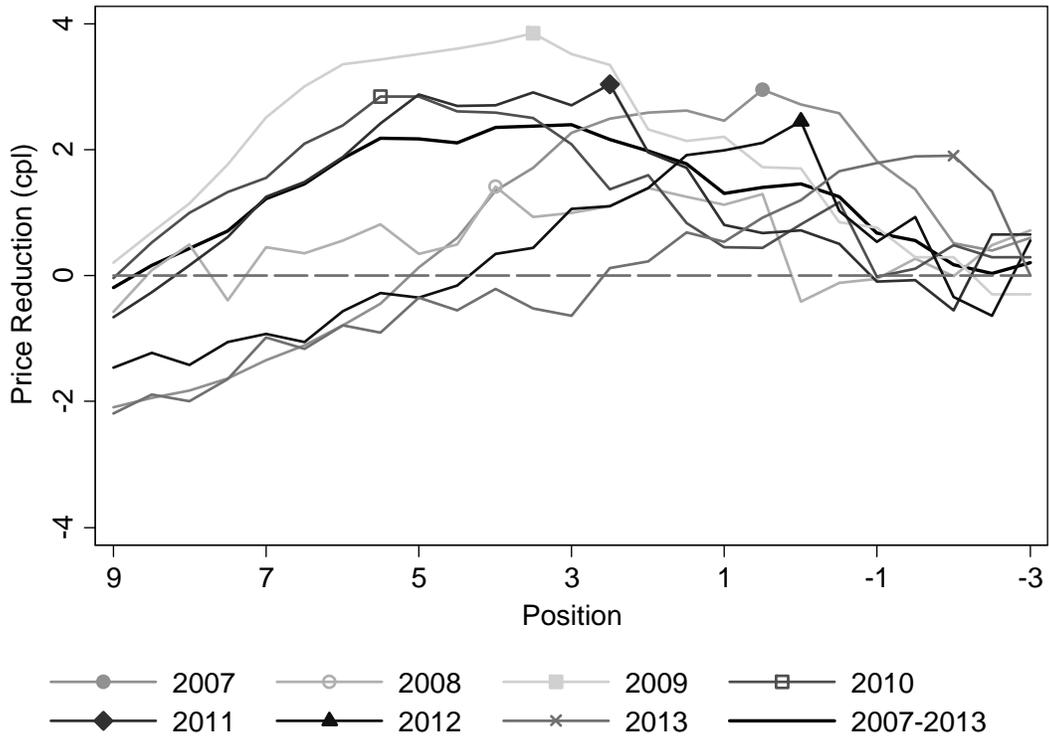


Figure 4. Reduction Profile for the Spike and Drop Class, Melbourne

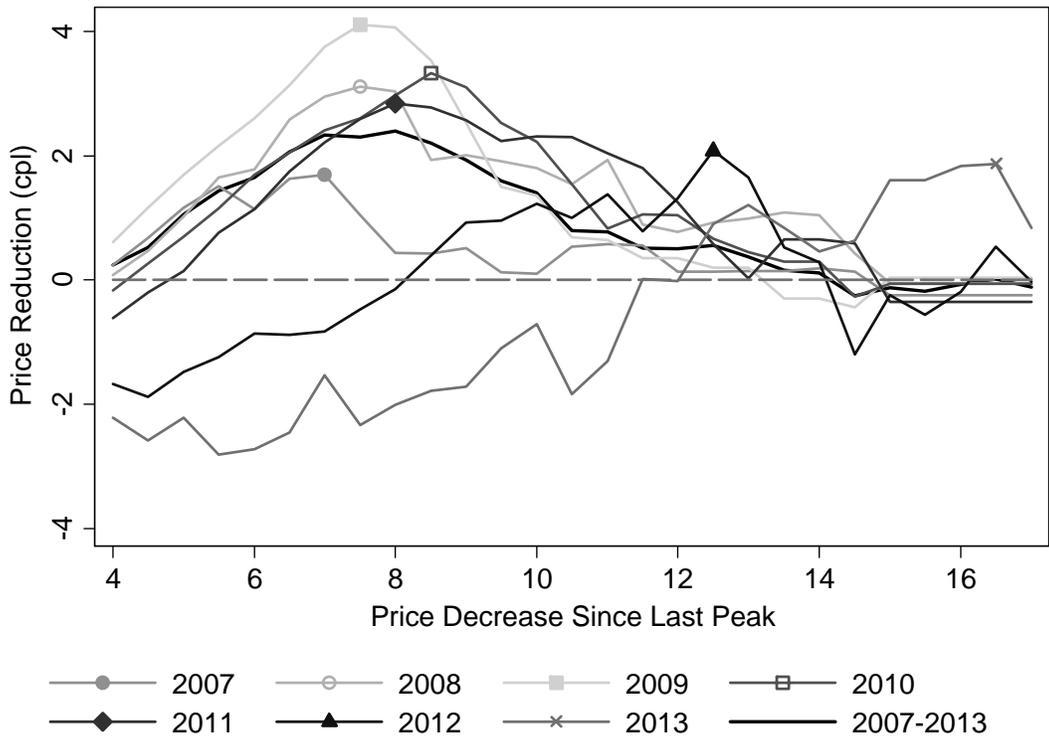


Figure 5. Reduction Profile for the Spike and Wait Class, Melbourne

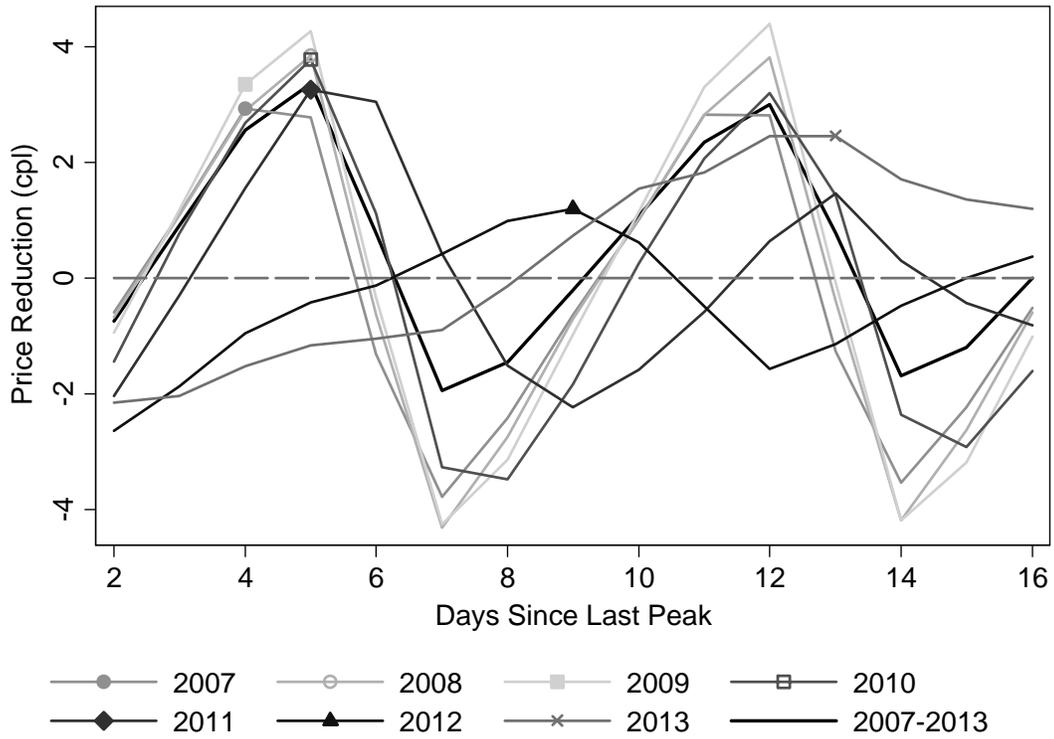


Figure 6. Reduction Profile for the Calendar-Based Class, Melbourne

